

Final Report

Evolvability of Multi-Agent Systems

A Collaborative research effort between the
Naval Research Laboratory and George Mason University
2004-2005

1. Introduction

The goal of this research was to analyze and improve the learning process for multi-agent systems using evolutionary algorithms (EAs). In particular, we wanted to take advantage of our previous work in developing a set of tools for analyzing the evolvability of genetic operators using Price's equation (Price 1970), a theory borrowed from the population genetics community.

All of our previous work with Price's equation (Potter, et.al. 2003; Bassett et.al. 2004) has been done on very traditional EA representations, specifically vectors of real values. Part of our research was to focus on how to adapt the use of Price's equation to a representation appropriate for controlling agents in a multi-agent environment. We chose to look at rule-based representations, in part because there are a large number of genetic operators defined for this representation that could be analyzed.

Another part of the research was to involve building multi-agent problem domains that could be used for testing our learning algorithms.

2. Price's Equation

Our main assumption going into this research was that the analysis tools we had developed based on Price's equation were at a point where they could be applied to any EA without difficulty. We soon discovered that this was not true.

Early in our research we observed anomalies in the results we received from our tool in certain circumstances. These results did not match well with other independent measurements that we made for verification purposes. This launched us down an avenue of research in an attempt to understand what was going on.

In particular, we noticed these problems when we used survival selection (when selection occurs after the genetic operators are applied) as opposed to parent selection (selection occurs before operators). As we did a detailed critical review of our own assumptions of Price's equation, we began to realize that we had not

20060404004

quite interpreted it correctly. There appears to be a built in assumption that selection occurs before the operators, at least if we want to keep interpreting the results the way we had been. So we began to consider what it would mean to interpret the results when survival selection is used.

It turns out that Price's equation still holds true when survival selection is used, but the terms of the equation take on different meanings. Put very simply, survival selection acts as a filtering process, weeding out the negative effects of the genetic operators so that only the positive effects are considered. This has some very useful properties. It allows us to get a better understanding of how much each operator is contributing to the evolutionary process. Thus we were able to improve on how our tool is used. These results were published (Bassett, et.al. 2005), and in the paper we demonstrate that an EA can be viewed as using parent selection or survival selection depending on ones "frame of reference" with regards to intermediate populations. This means that this new approach can be generalized to any EA.

Because of this detour, and time spent working on problem domains, we have not yet had a chance to start applying our tool to new representations, but we hope to begin that soon.

3. Multi-Agent Problem Domains

A fair amount of our time has been spent building multi-agent problem domains. We have constructed a number of these domains using the Player/Stage package, including a herding domain, an evasion-and-capture domain, and a harbor defense domain. Initial work was also done on a generalized capture-the-flag domain.

Player/Stage has presented us with a number of problems to overcome. The biggest one now is the computational resources required to do a learning experiment. Unfortunately, the simulation can only run at real-time speeds, making it difficult to perform learning experiments in a reasonable amount of time. The developers of Player/Stage have suggested that simulation speeds could be significantly improved by bypassing the Player module completely and making calls directly to the Stage library (the simulator). Initial tests with this approach have shown some success, but there are questions about whether this is the best approach in the long run. A significant amount of work will need to be done in order to write to these new APIs, and in a package under constant development, the way Player/Stage is, these APIs are likely to change in the future.

Another approach we are considering is to use a light-weight simulator (like MASON) for our initial experiments. Then, using transfer learning, we can bring the results of those experiments over to the Player/Stage environment to finalize

the process. Towards this end we have already built a harbor defense domain using MASON. Hopefully we can begin some transfer learning experiments soon to test whether this approach will be useful or not.

4. Conclusions

There are two main results of our collaborative effort. The first is significant improvements to our tool which uses Price's equation for analyzing EA dynamics. The second is continuing improvements to our suite of multi-agent problem domains. These contributions will benefit both NRL and GMU.

5. References

Price, G. (1970) Selection and Covariance. *Nature* 227 pp. 520-521

Potter, M. A., Bassett, J. K. and De Jong, K. A. (2003) Visualizing Evolvability with Price's Equation. *Proceedings of the 2003 Congress on Evolutionary Computation*, pp. 2785--2790, IEEE

Bassett, J. K., Potter, M. A. and De Jong, K. A. (2004) Looking Under the EA Hood with Price's Equation. *Proceedings of the 2004 Genetic and Evolutionary Computation Conference (GECCO)*, Vol 1, pp. 913--922, Springer

Bassett, J. K., Potter, M. A. and De Jong, K. A. (2005) Comparing parent selection and survival selection using Price's equation. *Proceedings of the 2005 Genetic and Evolutionary Computation Conference (GECCO)*, Springer